Los Alamos
NATIONAL LABORATORY
— EST. 1943 —

LA-UR-21-24951

Title: Maintenance Engineering Basics & Best Practices

Author(s): Sherwood, John Gregory

Intended for: Report

Issued: 2021-05-21

Maintenance Engineering Basics & Best Practices

John Sherwood, E-2

May 2021

Table of Contents

This page intentionally has no information

Introduction

The purpose of this document is to concisely describe the basics of maintenance engineering, summarize maintenance policies at Los Alamos National Laboratory, and also review current industry best-practices through the use of case studies. Maintenance engineering incorporates the study of failure, reliability, policies, and procedures in order to optimize cost and equipment availability. Each of these topics will be covered within this report.

1. Background

Novelist Chinua Achebe evoked a surprising depth from the simple title of his magnum opus, *Things Fall Apart*. The adage is certainly true for machinery and equipment – nearly every object degrades over its lifetime and eventually fails. Maintenance attempts to compensate for this degradation through cleaning, adjustments, recalibration, or refurbishment. And while people have been performing maintenance since tools have existed, maintenance engineering is a relatively new discipline. Maintenance engineers apply engineering concepts to optimize maintenance procedures, policies, and budgets in order to maximize equipment reliability,[1] maintainability,[2] or availability[3] [1].

1.1. History of maintenance engineering

Prior to the 1950s, maintenance was largely corrective. People tended to follow the saying, "if it ain't broke, don't fix it." As such, few tried to predict, plan for, and mitigate future failure events through preventive maintenance actions; these were considered a waste of resources. Corrective maintenance was merely an unavoidable cost.

However, after WWII, the growing adoption of operations research practices also became the advent of modern maintenance engineering. That is, companies began to see optimization opportunities in maintenance activities to reduce costs of corrective maintenance through preventive actions. Companies began to predict machine lifetimes and saw preventive maintenance as a long-term cost-savings measure.

The 1960s and 70s saw more shifts in maintenance thinking spurred by several independent catalysts, all related to the growing complexity of equipment and machines. One catalyst was the Dept. of Defense beginning to recommend or require a life-cycle costing approach to new large-scale procurements. As maintenance is part of a machine's life-cycle cost, maintenance costs needed to be predicted with accuracy. In parallel, United Airlines, competitors, and the Federal

---

[1] **Reliability** – the conditional probability that a system, product, or service within a given physical operating condition will successfully complete a mission of specified duration in a prescribed operating environment without disruption [9].

[2] **Maintainability** – the relative ease, in time and resources, with which an item can be retained in, or restored to, a specified condition when maintenance is performed. It is a function of design [24].

[3] **Availability** – the probability that a system or product will be available on-demand to perform a mission [9].

Aviation Administration were developing Reliability-Centered Maintenance (RCM)[4] to maximize the reliability of airplanes while minimizing unnecessary preventive maintenance actions. Also in 70s, the Japan Institute of Plant Engineers developed the concept of Total Productive Maintenance (TPM),[5] which focuses on maintenance in a manufacturing context. TPM views maintenance in terms of impact on manufacturing processes, based on availability, production rate, and output quality. Effectively, both RCM and TPM led to new maintenance perspectives and alternative ways of generating maintenance policy.

The 1980s saw a trend towards Condition Based Maintenance (CBM)[6] with the use of sensors and instrumentation to predict failure and act accordingly. The increased ability to capture data on equipment health led to more advanced analytical models. Concurrently, the diffusion of electronics, computers, and digital record-keeping provided further opportunities to store, manage, and analyze maintenance data.

This trend has continued through today. While there has not been a significant paradigm shift since the rise of Condition Based Maintenance, incremental improvements have occurred alongside digitization – maintenance data from remote sensors can be wirelessly streamed to a database and analyzed within a Computerized Maintenance Management System (CMMS).[7] Some have termed this state-of-the-art as "e-Maintenance" (similar to e-commerce or e-mail); it may soon just become "maintenance" – that is, a standard without an effective alternate method for comparison.

The future will likely see the rise of internet-of-things (IoT) sensor technology playing a larger role. These sensors may even report back to equipment manufacturers. For example, Airbus has recently partnered with Palantir (a data analytics company) to stream equipment health data from all Airbus planes back to Airbus for further maintenance analysis (this data may also be shared among airlines using Airbus planes) [2]. Other companies, such as IBM, are developing next-gen CMMS software suites that incorporate machine-learning techniques that will take over much of the data analysis [3].

1.2. Report outline

This report begins with a summary of failure and the typical "bathtub curve" used to describe the failure of a population of items (Section 2). In Section 3, reliability is presented, including discussion of reliability distributions, metrics, and reliability of multi-component systems. Section

---

[4] **Reliability Centered Maintenance** – a method for determining maintenance requirements based on the analysis of functional failures and their consequences. RCM's purpose is not to keep every component from failing, rather, RCM focuses on minimizing the consequences of component failure while maintaining system function [9].

[5] **Total Productive Maintenance** – a framework for maintenance decisions that maximizes equipment utilization. TPM has been primarily used in manufacturing environments, where machine reliability is less important than maximizing assembly-line throughput [26].

[6] **Condition-Based Maintenance** – a maintenance strategy based on measuring the condition of equipment in order to assess when preventive maintenance actions are necessary for continued function [9].

[7] **Computerized Maintenance Management System** – an electronic database and application suite used for tracking equipment, their status, health, and related maintenance logs. The CMMS may also be used to analyze data, develop preventive maintenance policies, and can be a repository for maintenance requests and procedures.

4 introduces maintenance and summarizes maintenance actions, polices, and concepts. It also provides many common maintenance policies. Section 5 provides a brief summary of modeling maintenance policies and economic optimization. Section 6 discusses Computerized Maintenance Management Systems, including reasons why CMMS may fail in an organization. Section 7 provides further maintenance related reading at Los Alamos National Laboratory, including a list of current maintenance policies. In Section 8, three industry case studies are presented – each concludes with a key point for maintenance managers. Finally, Section 9 provides a brief outline for how to establish an effective maintenance program.

2. Basic failure

Maintenance attempts to maintain or extend the life of an item or system through intervention actions. These maintenance actions attempt to reduce or stall the inevitable failure of items. Failures[8] may occur through several mechanisms, though these mechanisms can be grouped into two categories: overstress and wear-out. Both of these categories can be explained through Figure 1.
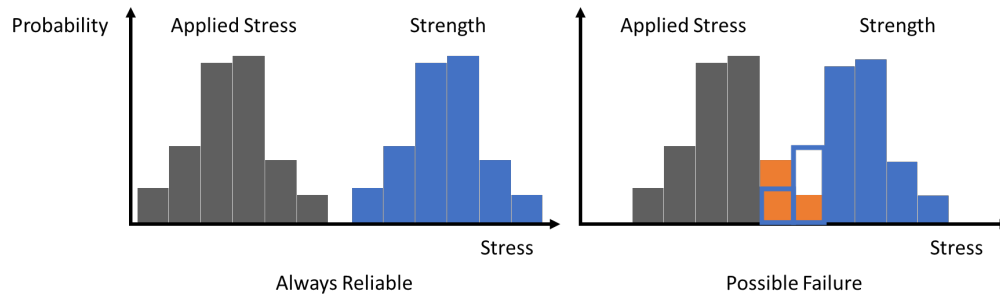

Figure 1. Failure occurs when an applied stress exceeds an item's strength. Adapted from Figure 3.1 of [4].

Figure 1 shows two graphs, each with a set of histograms. These histograms represent the probability of an environmental condition (colored grey), and the equipment of interest's probability of being able to handle the environmental condition (colored blue). For this figure, the performance metric is engineering stress. The equipment has an uncertain strength caused by some level of tolerance in its design. The local environment may also exhibit either some uncertainty or a range in the applied stress to the equipment.

Overstress occurs if the range of applied stress begins to overlap the equipment's strength. This results in a possibility of failure, shown as the orange bars in the right graphic. Here, the applied stress shifts to the right, indicating an increase of stress. Primarily, overstress is caused by misuse or mishandling of the equipment, or by changing environmental factors (e.g. temperature, humidity-level, air quality, etc.)

Wear-out is the inverse of overstress – rather than the applied stress shifting, the equipment's strength decreases. In a wear-out situation, the applied stress stays constant, but the equipment can no longer absorb it without failing. Primarily, wear-out is the result of aging, but may also be caused by a design or manufacturing failure (resulting in an inherent or induced weakness, below equipment specifications).

2.1. Types of failure

When failure occurs, it may either be immediate or gradual. That is, an item may fail through a two-state process, a many-state process, or an infinite-state process. These are represented in Figure 2.

---

[8] **Failure** – a condition in which an item or process can no longer fulfill its purpose or mission.
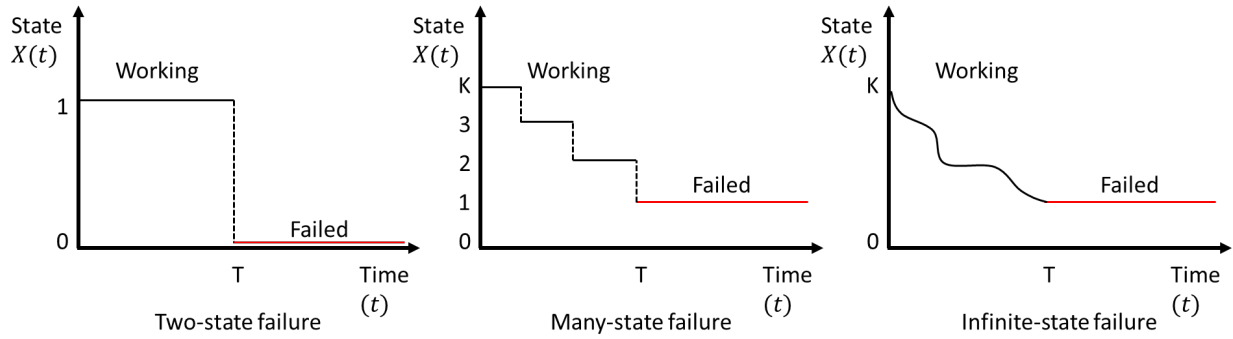
Figure 2. Comparison of failure mechanisms, including two-state, many-state, and infinite-state failure. Adapted from [5].

A two-state failure may be best represented by an incandescent lightbulb – it works exactly the same, right up until the point of failure (which is a complete failure). A many-state failure typically represents items with built-in redundancy, or situations where some components are not vital to the function of equipment. For example, a shopping cart wheel bearing may break, but the cart would still largely be usable. An infinite-state failure represents a continuous degradation of an item or component, such as engine oil, sandpaper, or the sole on a hiking boot. Note that, for many-state and infinite-state failure, there may be multiple failure states (with some being worse than others).

The states of an item may be enumerated and quantified for modeling purposes – each state may have a different probability of occurring, and the probability may be dependent on the state-history of the item. One way of modeling item states is through a Markov-chain model, though state-history adds a layer of complexity to the modeling process [6].

2.2. Failure severity

As mentioned above, there may be multiple failure states of varying severity. While there are multiple definitions or rankings of failure severity levels, a common classification system is found in risk management and system engineering. Both the Department of Defense and the Department of Energy have standards or guides describing the severity or consequences of risk, which can be applied to equipment failure. These classifications are shown in Table 1.

5

Table 1. A comparison of DOD and DOE risk severity categories.

Classification	DOD Criteria [7]	DOE Criteria [8]
Catastrophic (DOE: Crisis)	One or more of: Death, irreversible environmental damage, significantly large monetary loss	Project stopped. Withdrawal of scope or funding, severe schedule and cost performance issues.
Critical	One or more of: Permanent disability, hospitalization of 3+ personnel, significant (reversible) environmental damage, large monetary loss	Goals and objectives are not achievable. Additional funding and time may be required. Missed regulatory milestones and financial penalties.
Significant	N/A	Significant degradation in meeting objectives, significant increase in cost and schedule slip
Marginal	One or more of: injury resulting in 1+ lost work day, moderate environmental damage, moderate monetary loss	Small degradation in meeting objectives with marginal schedule and cost impacts.
Negligible	One or more of: minor injury, minor environmental damage, minor monetary loss	Minimal or no consequence to cost or schedule.

In both cases, the severity or consequence of failure is defined in economic terms: *how much will this cost the organization?* Additional important metrics include schedule delays, safety, and environmental impact.

From a maintenance policy perspective, the severity of failure is a key metric for determining maintenance actions. More severe failures ought to take priority over less severe failures, all else being equal. Of course, other factors (such as probability of failure and ease-of-maintenance) may affect decisions as well.

2.3. The bathtub curve

Another aspect of failure is failures across a population of the same items. Intuitively, identical machines should likely fail in similar manners, and at similar times. One useful conceptual model of this is the bathtub curve, shown in Figure 3. This curve is characterized by three distinct sections: 1) a decreasing failure region, 2) a stabilized failure region, and 3) an increasing failure region. The decreasing failure region occurs first and represents failures caused by inherent defects, such as manufacturing errors. Products that fail within this region likely never met their specification, and as such go through a "weed-out" process when put into operation. After this weed-out occurs, the remaining products presumably meet specifications and fail through random chance in the stabilized failure region. Finally, as products near their end-of-life, the probability of failure begins to increase through wear-out.
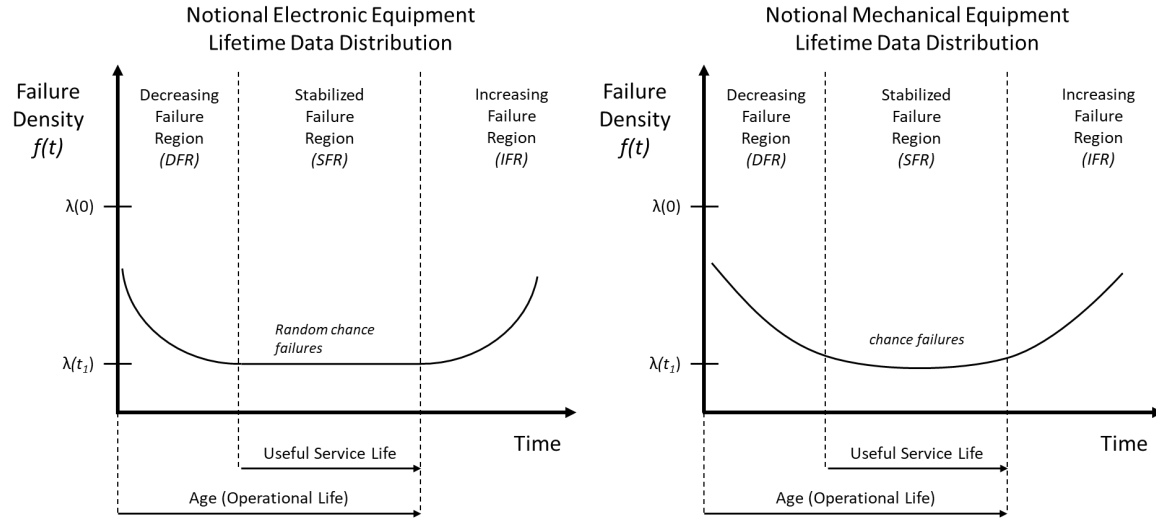
Figure 3. Conceptual model of failures across a population of either electronic or mechanical equipment. Adapted from Figure 34.8 of [9].

Some authors distinguish different bathtub curves for electronic and mechanical equipment. Electronic equipment tends to be more solid-state with few, if any, moving parts. As such, the stabilized failure region for electronics will be flat – closer to random chance of failure – compared to mechanical equipment that nearly always experiences minor wear. A flat failure region tends to be useful for modeling, as it can be simplified to a constant hazard rate[9] through time [9]. Because of the ease of modeling, some modelers simplify a mechanical system's stabilized failure region into a flat hazard rate.

Note that, despite the usefulness and conceptual intuitiveness of the bathtub curve, some researchers challenge the validity of the model [10] [11]. Smith notes that "the Bathtub Curve has been assumed to be applicable to more components than is supported by actual field data measurements" [12]. If data on a specific machine or component is available, it is better to construct a tailored model for it than to rely on a more generic model like the bathtub curve. Nevertheless, the bathtub curve is useful in developing a conceptual understanding of failure.

---

[9] **Hazard Rate** – the ratio of a system's failure rate to its reliability, as a function of time. $h(t) = f(t)/R(t)$

3. Reliability

3.1. Reliability & failure
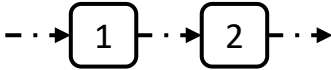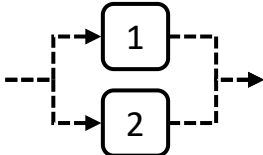
Failure and reliability are often contrasted. Reliability is the probability that a machine will not fail within a given timeframe (e.g. the time required to complete its mission). That is, reliability is a function of time, and reliability $R(t) = 1 - P_{failure}(t)$. If failure is modeled using a constant hazard rate (such as assuming a stabilized failure region of the bathtub curve), this equation becomes $R(t) = 1 - \int_0^\infty \lambda e^{-\lambda t} dt$, where λ is the hazard rate [9]. This hazard rate is usually expressed as the number of failures per time period, such as 0.3 failures per month.

The reliability of a system generally decreases over time – if the hazard rate of a machine is 0.5 per hour, each hour is effectively a coin-flip to determine if the machine will fail. Over time, the probability of these successive coin flips will eventually catch up with the machine, resulting in a failure.

3.2. Reliability & failure of multiple component systems

Reliability can be modeled for a single piece of equipment, or for multiple pieces of equipment within a system. For equipment in series or in parallel, the rules of calculating system reliability are the same as with resistances in electrical circuits. These equations are shown in Table 2.

Table 2. System reliability equations for two-component configurations.

Configuration	Diagram	Equation
Series		$R_{sys}(t) = R_1(t)R_2(t)$
Parallel		$R_{sys}(t) = R_1(t) + R_2(t) - R_1(t)R_2(t)$

Running equipment in parallel (or series) increases (or decreases) system reliability. However, this additional redundancy (or additional risk) is only apparent for a specific time period. The time period is dependent on the underlying failure model, though it is generally the time before failure of a single component becomes very likely (e.g. before reliability reaches 10%).[10] This is demonstrated in Figure 4.

---

[10] Other failure models, such as one of increasing failure over time, show minimal apparent difference in system reliability near the start of system operational time (e.g. when reliability of a single component is above 90%).
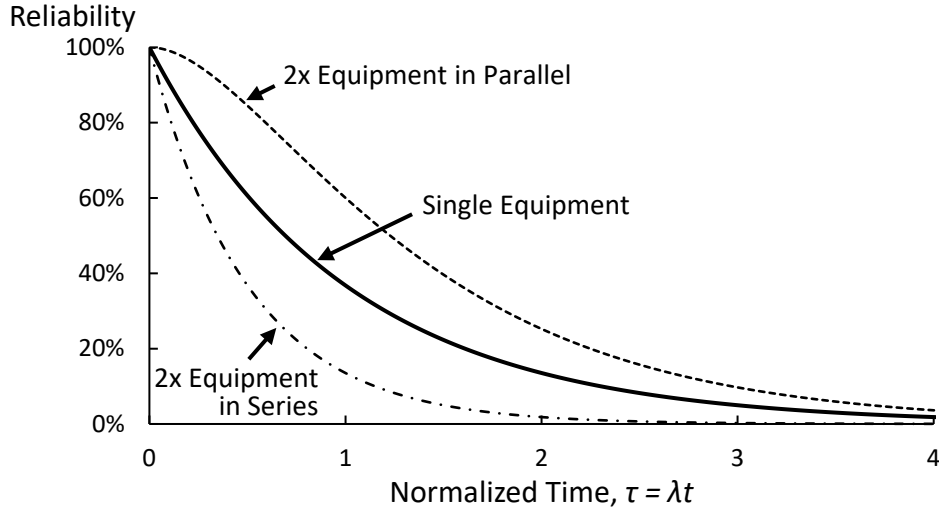
Figure 4. Reliability over normalized time using a constant hazard rate, adapted from [13].

Figure 4 shows the reliability of simple systems as a function of normalized time. Here, normalized time is equivalent to the hazard rate multiplied by some time unit. That is, a normalized time of $1 = 0.1 \left[\frac{failures}{hour}\right] \times 10 \ [hour]$ or $1 = 0.01 \left[\frac{failures}{month}\right] \times 100 \ [month]$.[11] Using normalized time provides a simple means to convert a failure rate into the time dimension, in order to understand how reliability changes through time.

Using a constant hazard rate, Figure 4 shows that running equipment in parallel improves overall system reliability. The largest boost in reliability occurs at a normalized time of 0.84, where reliability of the parallel system is 24.6% better than using a single piece of equipment in isolation. After this point, the benefits of a parallel system begin to slowly taper off.

3.3. Reliability distributions

The above discussion focuses only on the case of a constant hazard rate. Lee cites the constant hazard rate, or exponential distribution, as one of five useful probability distributions for describing failure and reliability [13]. The five useful distributions are described in Table 3.

---

[11] Note that, despite the units of normalized time being [failures], normalized time does not capture the probability of failure. A constant failure rate of 0.01 per month does not imply failure will occur at (or before) month #100. Rather, (in the case of a constant failure rate), this represents the mean time to failure.

Table 3. Useful probability distributions for describing failure, adapted from [13].

Distribution	Description	Used for…
Exponential	Represents a constant failure rate – failure has an equal probability of occurring through time	Characterizing random chance failures
Erlangian	Represents multiple exponential functions "stacked" on top of each other, the number of which is represented by a k parameter	Characterizing a system with independent, identical components each facing a random chance of failure. System failure occurs after the k^{th} component fails
Gamma	Exponential and Erlangian are special cases of the Gamma function; Gamma allows for more flexibility to fit data	Characterizing fatigue failures arising from repetitive shocks
Lognormal	Has the form of a normal distribution, but disallows negative values	Useful when probability of events may occur over several orders of magnitude
Weibull	Includes both exponential and normal distributions as special cases	Widely used because it encompasses all cases in which the hazard rate varies according to some power of t, time

The choice of probability distribution (and therefore form of hazard rate) is of critical importance to accurately model reliability or failure. Reliability engineers often collect failure data which are used to fit models; the best-fitting model is used for further analysis. Typically, models are fit using maximum likelihood estimation rather than a least-squares approach [14]. Note that inaccurate data or too small a sample size can have outsized effects on the model accuracy, and therefore estimates of equipment reliability.

If there are limited data, a Bayesian approach to modeling is recommended [15]. A Bayesian approach allows for more flexibility to incorporate additional failure data as it becomes available (compared to a classical statistical model). Bayesian statistics incorporates prior knowledge (the "prior distribution") alongside current observations to generate an up-to-date model (the "posterior distribution") [16].

3.4. Reliability metrics

Models of reliability are used to determine various metrics related to equipment lifetimes. A key metric is the mean time to failure (MTTF). Mathematically, $MTTF = \int_0^\infty R(t)dt$. In the case of an exponential distribution with constant hazard rate, $MTTF = 1/\lambda$. A related metric is the mean time between failure (MTBF). MTBF includes repair times, such that $MTBF = MTTF + MTTR$, where MTTR is the mean time to repair.[12]

---

[12] **Mean Time To Repair (MTTR)** – the average length of time needed for repair. Typically, this is estimated based on historical maintenance times and is often considered "0" if the repair time is inconsequential relative to the MTTF.

These metrics are used to determine equipment availability.[13] Availability represents the probability that a system will be operational when needed. It is calculated as total availability $A_\infty = MTTF/MTBF$, which is the ratio of machine uptime to total time (uptime + downtime) [17].

Availability and reliability are occasionally conflated. Reliability is the probability that a machine will work for a period of time *without disruption*, whereas availability is the fraction of time that a system is operational. Stakeholders may state a desire to know the reliability of something when they may actually care about availability.

While reliability is important to maintenance, reliability engineering as a discipline seems to be focused on ways to improve the reliability of products under development – *how can we design this to be more reliable?* Reliability engineers may also be involved in generating reliability requirements and policies such as warranty offerings. As this report focuses on maintenance engineering, we turn to maintenance.

---

[13] **Availability** – the probability that a system, product, or service will be available on-demand to perform a mission.

4. Maintenance

Things are unreliable and fail. Maintenance, broadly speaking, is a framework used to keep equipment operational. Most commonly associated with maintenance are maintenance actions – repairs or interventions carried out on a component or device. Maintenance also encompasses maintenance policies. In this context, maintenance policies describe the triggering mechanism for various actions. Even broader than policies are what Ben-Daya et al. call the maintenance concept [18]. A maintenance concept is the set of policies and also general decision structure in which policies are developed an enacted. In short, the maintenance concept is an institution's strategic perspective on maintenance. These three levels of maintenance (action, policy, and concept) are all crucial for effective maintenance. We will discuss each in turn.

4.1. Maintenance actions

Maintenance actions are direct interventions in a system, typically carried out by technicians. Actions are not just limited to repair; actions may include inspections, compliance testing, and monitoring in addition to repair work. Actions are typically classified as corrective or preventive. Corrective maintenance encompasses interventions that occur as, or after, a failure arises. Preventive maintenance is defined as actions that occur according to a prescribe criteria of time, usage, or condition, and is intended to reduce the probability of failure prior to the failure occurring.

4.1.1. Maintenance classification

Preventive maintenance may be predetermined, condition-based, or opportunistic. Predetermined maintenance is usually specified through a maintenance policy, and may be based on various "clocks." Three common "clocks" are:

1) Calendar clock: maintenance occurs based on time of year or season. Repair or replacement occurs regardless of usage. An example may be changing to summer or winter tires on a car.
2) Age clock: maintenance occurs based on lifetime of equipment. Repair occurs after a number of hours, regardless of usage. This is common with materials that naturally break down over time, like rubber, various chemicals, or UV sensitive plastics.
3) Usage clock: maintenance occurs based on equipment runtime. Repair occurs after a number of operating hours (or other relevant usage metric), such as replacing printer ink after 300 prints.

Age and usage clocks are related to each other. A piece of equipment cannot be used for more hours than it is old. Many products and equipment use both clocks to inform maintenance decisions – car engines "need" oil changed every 3,000 miles or every three months. Graphically, this can be shown as Figure 5.
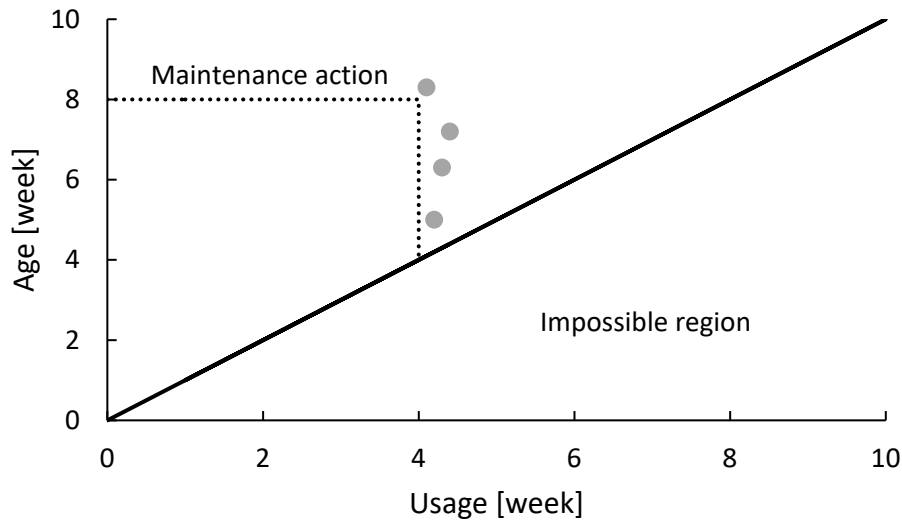
Figure 5. Comparison of usage clock and age clock. Usage cannot be higher than age. An example preventive maintenance policy is shown as a dotted line. The failure data points demonstrate an instance where a usage policy would be effective.

Here, a preventive maintenance policy that incorporates both age and usage is shown as a dotted line. If equipment passes that line (either through excessive age or excessive usage), preventive maintenance would be enacted. The graph also shows failure data for the equipment. Because this failure data mainly occurs as a function of usage rather than age, a usage-clock maintenance policy alone would be effective.

Condition-based maintenance[14] is another type of maintenance action. Conditioned-based maintenance involves equipment monitoring through either periodic inspection or sensor data. Under a condition-based framework, maintenance policies tend to be more individualized. Various metrics are monitored for specific components and, once a metric hits a threshold, preventive maintenance will be enacted. An example metric is monitoring particle debris in oil or lubricants. Monitoring this metric allows for trend analysis to help predict when failure might occur. Because of this increased predictive capacity, condition-based maintenance is a preferred approach.

The final type of preventive maintenance is opportunistic maintenance. Opportunistic maintenance occurs if a part is easily replaceable during the maintenance of another component. So, the component may receive maintenance earlier than required, but there are cost-savings associated with performing maintenance on several components at once. Typically, opportunistic maintenance is difficult to plan. The maintenance technician or manager are usually best-equipped to determine opportunistic maintenance. As such, an institution's maintenance concept ought to enable opportunistic maintenance decision-making at the boots-on-the-ground level.

An overview of these maintenance classifications is provided in Figure 6.

[14] Condition-based maintenance is also known as predictive maintenance.
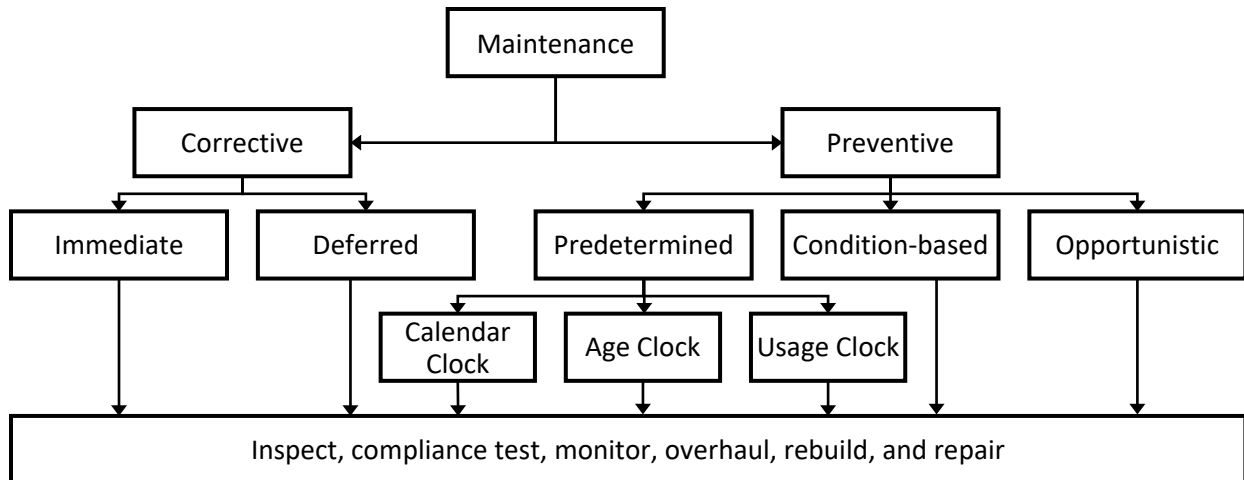
Figure 6. An overview of maintenance classifications. Adapted from Figure 4.1 of [18]

4.1.2. Repair

Items that fail can be repaired. Not all repairs are equal, however. A repair can have two types of effects on a reliability function. Recall that reliability varies over time – a repair can "reset" the reliability clock, shifting the reliability function backwards in time. If identical repair parts are used, typically only the reliability clock shifts. Alternatively, a repair can alter the machine's hazard rate – this changes the reliability function itself. The reliability function will change if a repair part is better or worse than the original. Ben-Daya et al. define four levels of repair [18]. The four different levels delineate the overall effectiveness of a repair. These four levels are shown in Figure 7.
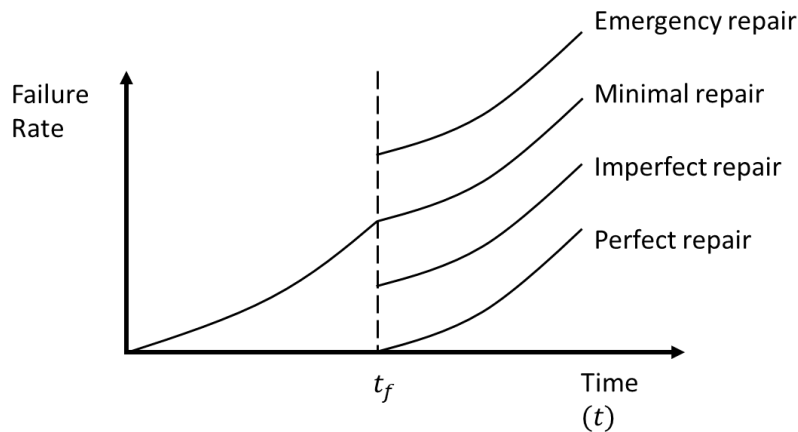
Figure 7. Four levels of repair and their effect on failure rate, adapted from Figure 4.4 of [18].

Emergency repair is best described as a duct-tape fix. After a failure and emergency repair, the reliability becomes worse than prior to failure. Typically, an emergency repair alters the reliability function (rather than running the reliability clock forward).

Minimal repair is a type of repair where the reliability fundamentally does not change. The repair process, outside of fixing the failed equipment, has no impact on the failure rate or reliability clock.

14

The effect of this repair is as if no failure occurred, but everything else about the equipment remains unchanged.

Imperfect repair is a fairly common repair type. Within imperfect repair, the reliability clock is wound back but not to brand-new. Imperfect repair can occur when repairing with used parts, but it more often occurs when only a subset of parts are replaced. If an assembly fails due to a single part breaking, that part may be replaced. But, other components of that assembly still have wear. Therefore, the reliability is diminished relative to a brand-new assembly.

A perfect repair is a complete replacement of damaged or degraded components, such that the equipment is "good-as-new." This effectively resets the reliability clock and the machine or component is treated as if it has no operational time on its usage clock.

It is usually the case that perfect repairs are the most expensive to enact, followed by imperfect, minimal, then emergency. Therefore, there is a cost/reliability tradeoff to consider. There is also room for optimization.

4.2. Maintenance policy

A maintenance policy is a set of rules describing the triggering mechanism of maintenance actions. Policies may be a single rule (e.g. "change oil after 3,000 miles") or more complex (If A occurs then take action 1, else if B occurs then take action 2). Note that this is a narrow definition of policy – within this context, a policy is not an institutional directive providing guidance at a broader scale. Ben-Daya et al. describe several basic maintenance policies, each discussed below [18].

4.2.1. Clock-based policy

In an age-based policy, components are replaced when they reach a predetermined age or, are replaced if they fail prematurely. This policy includes both preventive maintenance (PM) and corrective maintenance (CM) actions. A policymaker must choose the age of replacement, along with which type of clock to use.

4.2.2. Block policy

Similar to a clock-based policy, a block policy specifies replacement according to clock intervals (e.g. every three months). Block policies tend to be used for items required frequent maintenance (fluids, seals, limited-life parts), whereas a clock-based policy may apply to items requiring infrequent maintenance (e.g. engine or transmission replacement). As with a clock-based policy, premature failures are repaired as they occur and reset the interval. A policymaker must choose the clock interval.

4.2.3. Periodic policy

A periodic policy also specifies replacements according to a clock interval, but either ignores or minimally repairs premature failures. Corrective maintenance does not reset the interval in a periodic policy. An example may be lights in a warehouse or high-bay production environment – if one light fails, it is costly to stop production or rent a lift. Because a single light failing is not necessarily a safety concern, the maintenance team may choose to defer repair until the periodic

policy calls for all lights to be replaced. A policymaker must choose the clock interval and ensure premature failures are low impact (either through a low consequence or low likelihood, or both.)

4.2.4. Failure limit policy

A failure limit policy replaces equipment once its reliability or failure rate reaches a certain threshold. As a reminder, reliability decreases over time as failure becomes more likely. Before this reliability threshold is reached, any failures are repaired through minimal or imperfect repair. An example of this policy may be the average car – people tend to perform imperfect repairs (i.e. not an overhaul) until the car is deemed "unreliable" from imminent cascading failures. A policymaker needs to choose a reliability or failure rate threshold to enact this policy.

4.2.5. Repair cost limit policy

This policy type takes into account the cost of repairs. A pure version of this policy does not use preventive maintenance. Rather, repair costs are estimated at the time of failure. If the estimated repair costs are more expensive than either a set limit, or the cost of replacement, then the equipment is replaced. This policy may be somewhat dynamic because replacement costs may change over time. A policymaker may choose to do an estimate at the time of policy enactment to see if this policy is worthwhile. The policymaker may also choose a static threshold (e.g. $50,000 or 5% of management reserve) if desired.

4.2.6. Repair time limit policy

Similar to a repair cost limit policy, a time limit policy requires estimating the time required for a repair. If the repair takes longer than acceptable (resulting in a drop of availability), then the item is replaced instead. Typically, this policy is used when the cost of time lost is much greater than either the repair or replacement price. This policy may be useful to help prevent schedule slippage, particularly if replacement equipment is readily available.

4.2.7. Repair count policy

The final basic policy is a repair count policy. For this type of policy, maintenance managers track the number of times a repair has occurred. After the Kth failure, the unit is replaced rather than repaired. This policy has no direct preventive maintenance actions, it merely tracked corrective maintenance. The repair count policy is designed to counteract ever-increasing failure rates caused by repeated imperfect repair. For many objects, repair cannot be perfect because of general overall degradation (e.g. embrittlement, micro-fracture propagation, etc.)

4.2.8. Important policy considerations

The choice of policy (or policies) to implement is dependent on a few different dimensions. These include things like cost, repair time, reliability, and others. A full maintenance policy should address these dimensions to best describe and fully specify the conditions and requirements for maintenance actions. A list of relevant dimensions is provided in Table 4. The table also provides two example objects with which to craft a maintenance policy.

Table 4. List of relevant maintenance policy dimensions and their possible variables. Adapted from Table 4.2 of [18]

Dimension	Potential Variables	Example: light bulb	Example: car engine
Item	Single or multi component system	Single	Multi
Types of failure	Two-level, discrete, infinite, intermittent	Two-level	Discrete
Monitoring plan	Continuous, intermittent	Intermittent	Continuous sensor data
Inspection interval	Planned (constant or variable), unplanned	Unplanned	Planned (3,000 miles)
Operating parameters	Known & specified (e.g. SPC chart), unknown	Known	Known w/ uncertainty
Time horizon	Finite (applicable for limited time), infinite	Infinite	Finite (until next major service)
Item usage	Continuous, intermittent	Intermittent	Intermittent
Level of repair	Emergency, minimal, imperfect, perfect	Replacement	Imperfect (used parts)
Repair time	Negligible, taken into account	Negligible	Take into account
Repair cost	Insignificant, significant	Insignificant	Significant
Replacement cost	Insignificant, significant	Insignificant	Significant

A maintenance policy for the light bulb may include the following. The light bulb's repair time and cost are minimal and insignificant. It will have unplanned inspections each time it is operated and, when it fails, it will be replaced. The light bulb may have an expected lifetime of 15,000 hours. As such, preventive maintenance will occur according to a clock-based policy when the bulb is seven years old (this assumes a standard usage based on work-hours per year.) If failure occurs prematurely, it will be replaced.

A maintenance policy for the car engine may be: the car engine is a multi-component system that may experience multiple discrete types of failure. It will be continuously monitored by the engine control unit which will activate the check-engine light and store a diagnostic code when an issue arises. The engine will be inspected for signs of wear during each oil change. If maintenance is required, it will be analyzed through a repair cost limit policy. If the maintenance costs more than the value of the car, consider the car totaled and replace. When possible, seek out used or aftermarket parts to minimize cost. This policy will be effective until the next major service milestone set at 100,000 miles. At that time, the maintenance policy will be reevaluated.

These examples are brief; maintenance policies can become quite in-depth. This is particularly true for multi-component systems where different components may need different procedures. Nevertheless, the examples show how the dimensions in Table 4 help inform policy development.

4.3. Maintenance Concept

Maintenance concept describes the strategic stance that an institution takes regarding maintenance. This may involve how a maintenance department (or lack of departments, if maintenance staff are

integrated into other teams) is structured. Maintenance concept also describes how maintenance policies are generated, who has the authority to do maintenance work, and who is able to make maintenance decisions. Ben-Daya et al. states:

> "Effective maintenance decisions need to be made in a framework that takes into account technical, commercial, social, and managerial issues from an overall business perspective… Therefore, effective maintenance management needs to be based on quantitative business models that integrate maintenance with other decisions, such as those involving production, and so on." [18]

Ultimately, a maintenance concept outlines and prioritizes maintenance objectives, and creates a framework within the institution to meet those objectives.

An institution will usually have long-term objectives of productivity and cost-management. Short-term objectives must be aligned with long-term goals and established to ensure that the maintenance business is on the right track in fulfilling its mission. Some short-term objectives may include:

- Overtime reduction by some specified percentage;
- Implementation of a PM program for a specific area or group of equipment;
- Implementation of a work order system within a given area;
- Implementation of an inventory management program.

As an example, The Los Alamos National Laboratory Detonator Production Division is in the process of implementing an asset management program with the intent to better implement preventive maintenance. This will help meet a long-term objective of productivity required for mission needs.

Maintenance objectives tend to be fulfilled within the context of a framework or outlook on maintenance. Two common frameworks (mentioned in Section 1.1) are Reliability-Centered Maintenance (RCM) and Total Productive Maintenance (TPM).

4.3.1. Reliability-Centered Maintenance
The focus of Reliability-Centered Maintenance is on maintaining system function, rather than keeping a system in perfect condition. The failure of non-essential parts is inconsequential, and stopping work for these repairs provides no benefit. Is it worth fixing a broken radio in a fleet vehicle? RCM would say no.

RCM has four key steps:

1. Identify essential system functions and their mechanisms,[15]
2. Identify system function failure modes,

---

[15] Moubray and Ben-Daya et al. identify questions for subject-matter experts to aid in this process [25]:
    1. What are the functions and associated performance standards of the asset in its present operating context?
    2. In what ways can the asset fail to fulfill its functions?
    3. What causes each functional failure?
    4. What happens when each failure occurs?
    5. In what way does each failure matter?
    6. What can be done to predict or prevent each failure?
    7. What should be done if a suitable proactive task cannot be found?

3. Prioritize these failure modes, and
4. Select maintenance policies to mitigate these failure modes.

4.3.2. Total Productive Maintenance

Total Productive Maintenance arose from manufacturing environments and is well suited for production activities [19]. TPM is focused on overall equipment effectiveness, which includes performance and quality in addition to equipment availability. A traditional TPM framework has the following features:

- Equipment effectiveness (defined below) is maximized.
- A thorough policy set of PM actions is established for the equipment's entire life span.
- It involves every employee, from management to technicians within and outside of maintenance.
- PM is promoted through motivation and grass-roots employee empowerment.

The TPM process starts by identifying the major losses with regard to equipment. The following six losses limit equipment effectiveness (also shown on Figure 8):

1. Equipment failure (breakdown);
2. Set-up and adjustment downtime;
3. Idling and minor stoppages;
4. Reduced speed;
5. Process defects;
6. Reduced yield.

These six losses are linked to three metrics that affect equipment effectiveness: availability, performance, and quality. Overall equipment effectiveness (OEE) is defined as the product these three metrics.

$$A = \frac{\text{Operating time}}{\text{Loading time}}$$

$$P = \frac{\text{Widgets produced} \times \text{Ideal cycle time}}{\text{Operating time}}$$

$$Q = \frac{\text{Widgets produced} - \text{rejects}}{\text{Widgets produced}}$$

Overall equipment effectiveness (OEE)

$$OEE = A \times P \times Q$$
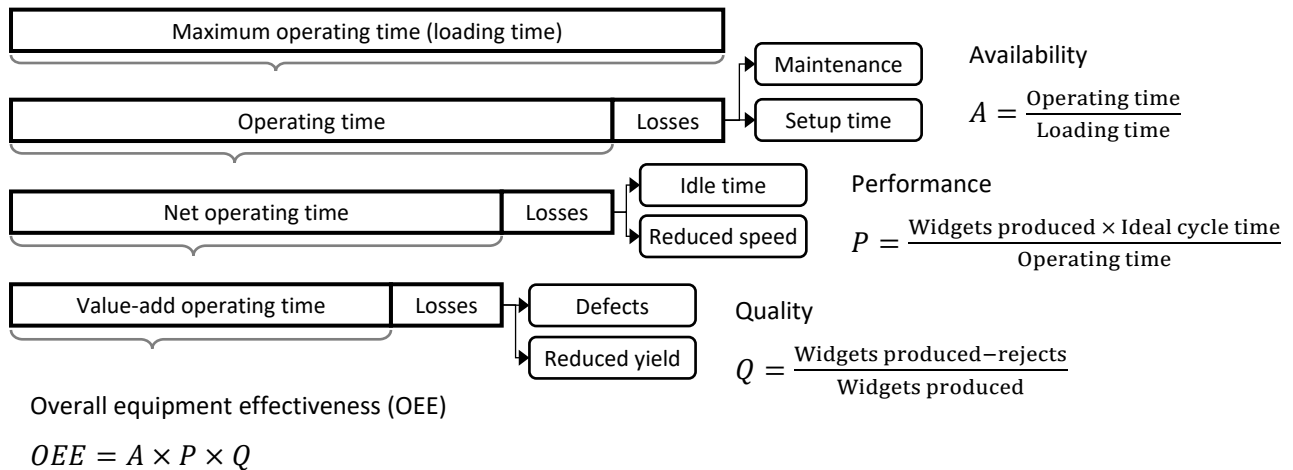
Figure 8. A graphical overview of Total Production Maintenance. The six major losses are shown, along with the equation used to calculate overall equipment effectiveness. Adapted from Figure 17.9 of [20].

For reference, an article in the trade magazine *Reliable Plant* states that a typical OEE is about 60%. An OEE of 40% is low but not uncommon, while an OEE of 85% is considered world-class [21].

Implementing TPM can occur in a narrow sense – maintenance managers can track OEE and improve it where possible. Or, TPM can be implemented across an organization, similar to continuous improvement programs. Note that TPM may seem similar to Total Quality Improvement (TQM); the difference is a focus on improving losses rather than a focus on improving quality.

5. Maintenance modeling & economics

Maintenance modeling is concerned with predicting and understanding the effects of maintenance policies and actions on different performance indicators. These indicators may be constrained to the technical level, including reliability, availability, and overall equipment effectiveness. Or, the indicators may be economic in nature: cost rate, expenditures, profits, etc. The performance indicators may also include other metrics, such as safety, quality, and risk.

Typically, a modeler will choose an indicator to optimize through an objective function (or vector of functions, if optimizing multiple indicators). This function may take the form:

$$J = \frac{ECC}{ECL} = \frac{Expected\ Total\ Cost\ per\ Cycle}{Expected\ Cycle\ Length} \tag{1}$$

where J is the objective function to be optimized. In general, J represents a total cost per cycle length – it is normalized to the mean time between failures. (J could also be normalized on a per fiscal year or other basis, but this can increase model complicatedness.)

5.1. Corrective maintenance only

In the case of an only corrective maintenance policy with inconsequential repair time, the objective function may become:

$$J = \frac{ECC}{ECL} = \frac{Expected\ Cost}{MTTF} = \frac{C}{\int_0^\infty R(t)dt} \tag{2}$$

Here, the expected cycle length reduces to the MTTF because repair times are minimal.[16] The maintenance cost may consist of a few items: $C = C_{direct} + C_{indirect} + C_{consequential}$

- *Direct Costs* (C_{direct}) are material and direct labor costs
- *Indirect Costs* ($C_{indirect}$) are administrative costs, overhead, costs to hold spare parts, etc.
- *Consequential Costs* ($C_{consequential}$) include production losses stemming from downed equipment, penalties from delays, etc.

These costs may or may not apply to every situation. But, labor, administration, and repair parts are common expenses. Many of these costs come from estimates and may exhibit variability. To increase the accuracy of these cost estimations, historical records of similar repairs are critical.

These costs are often contrasted with a replacement cost – the cost of a new system. A repair or replace decision is often made. If the total repair costs are higher than the total replacement costs, the replacement option is optimal. Note also that, hidden within MTTF is the type of repair (minimal, imperfect, perfect, see Section 4.1.2). The type of repair influences future reliability, which affects future instances of the objective function. If repairs are not perfect, they will shorten

---

[16] Remember that MTBF = MTTF + MTTR, from Section 3.4

the MTTF in the next cycle. This may raise the per-cycle repair cost and shift a repair-or-replace decision.

A model similar to Equation (2) could be constructed to account for multiple repair cycles, though this requires data on how each repair affects a machine's reliability function.

5.2. Preventive maintenance modeling

To incorporate preventive maintenance into a model, a modeler must also account for the probability that failure occurs before preventive maintenance is enacted. A general objective function for an age-based preventive maintenance model is similar to the corrective maintenance model, but it is modified to account for the expected costs of corrective and preventive maintenance. For example:

$$J(T) = \frac{ECC}{ECL} = \frac{Expected\ Cost}{MTTF + MTTR} = \frac{C_p R(T) + C_{cr} F(T)}{\int_0^T R(t)dt + MTTR} \tag{3}$$

where $C_p R(T)$ represents the cost of preventive maintenance multiplied by the probability that the equipment survives until time T (the time of PM action). This is the expected cost of preventive maintenance. Similarly, $C_{cr} F(T)$ represents the cost of corrective maintenance multiplied by the probability of failure up to time T.

This model also includes MTTR because the repair time is large enough to take into account here. Repair time may include:

- Time to investigate faults
- Time needed to carry out repair
- Testing after repair
- Various logistics / admin times (e.g. lead time for repair parts)

It is not uncommon for lead times to be a significant component of repair times – spare parts would negate this lead time.

The goal of this model is to minimize cost by varying time T, which is when preventive maintenance occurs. If T is too short, PM may actually be unnecessary. If T is too long, CM may be needed too often, which might increase costs.

These optimization models can quickly become unwieldy for more detailed or advanced maintenance policies [22]. However, there are some general takeaways from these models. All rely heavily on (uncertain) reliability models and hazard rates. Uncertainty in reliability will make maintenance optimization quite challenging. But, other factors may also reduce maintenance costs. These include repair time, replacement time, and costs associated with repair. If reliability is uncertain, managers may choose to minimize these factors in order to reduce costs and increase availability.

6. Maintenance systems

Maintenance systems are computer or server-based application suites designed to handle the majority of maintenance engineering tasks. Most maintenance systems are called Computerized Maintenance Management Systems (CMMS), though they are sometimes referred to as Enterprise Asset Management (EAM) systems. These maintenance systems usually contain tools to track equipment and their maintenance history, maintenance plans and procedures, and work requests. They also may be used to analyze failure data. State-of-the-art systems also include capabilities like real-time sensor readings for use in condition-based maintenance monitoring. In general, a CMMS has three distinct elements or modules:

1) A database containing
   a. Asset register (i.e. a list of equipment)
   b. Asset conditions or current status
   c. Asset maintenance history
   d. Asset maintenance policies
   e. Preventive maintenance procedures library
   f. Maintenance personnel database, for use when assigning work orders
2) Statistical tools
   a. Models that can be applied to items in the asset register; the user can select an appropriate model and alter various parameters
   b. Reporting tools, such as dashboards, graphics, tables, and report generation
3) Decision Support System
   a. Scheduling tools for maintenance – frequency, procedure, assigning spare parts, assigning technicians
   b. Inventory management
   c. Metrics and KPIs to assess performance of a maintenance program or policy

These components, together, form a comprehensive system with which to manage maintenance. It can be very effective at increasing efficiency when properly used. There are a few reasons CMMS can fail, however. Note that the database must be accurate and up-to-date to be of any value – if technicians or managers treat the CMMS as a less-important, secondary system, the CMMS will quickly lose value and accuracy. (This may result in a reinforcing feedback loop!) CMMS may fail if there is inadequate training or buy-in from the staff interacting with it. Only well-trained personnel can use a CMMS to its full capability. Finally, CMMS may fail if the specific software suite selected does not adequately match institutional needs or barriers.

The future of CMMS and maintenance systems will see a shift towards what some call e-maintenance. These systems tend to be cloud or server-based and allow for direct connection to internet-based sensors. This level of connectivity increases the capabilities of condition-based maintenance models. Additionally, CMMS will see a shift towards enhanced statistical models that adapt over time: machine learning. CMMS may also become connected to product manufacturers to allow for sharing of maintenance data. This sharing would enhance the accuracy of predictive maintenance. Both IBM and Palantir have been demonstrating these concepts [2, 3].

7. Maintenance at Los Alamos National Laboratory

The Maintenance and Site Services (MSS) group within Facility & Operations is in charge of the Laboratory's institutional maintenance program. This program services both facility and programmatic equipment. However, the majority of available documents and policies seem to be focused on building and facility maintenance rather than programmatic equipment. As such, guidance on manufacturing equipment and other machinery may be limited or at the discretion of equipment owners. Table 5 lists maintenance policies currently at the Laboratory.

Table 5. Maintenance related policies and administrative procedures at Los Alamos National Laboratory

Policy	Name	Description
P950	Conduct of Maintenance	This document serves as a parent document to a series of implementing Administrative Procedures (APs) that detail maintenance and work management procedures at the Laboratory. It describes requirements for Preventive Maintenance (PM), Predictive Maintenance (PdM), and Corrective Maintenance (CM), and for the assessment and inspection of the condition of Structures, Systems, and Components (SSCs) during daily work routines and at designated frequencies. [23]
AP-MNT-002	Seasonal Facility Preservation	"Implements the requirements and the process for the development and implementation of seasonal facility preservation plans for Los Alamos National Laboratory (LANL) facilities to ensure continued safe facility operations." [24]
AP-MNT-003	Determining Maintenance Facility, Equipment, and Tool Needs	"Implements the requirements and process for the evaluation of maintenance facilities, equipment, and tools used in the performance of maintenance tasks at the Los Alamos National Laboratory (LANL). This procedure implements elements of P 950, Conduct of Maintenance." [25]
AP-MNT-004	Facility Condition Inspection	"Implements the requirements and processes for assessing the physical condition of Los Alamos National Laboratory (LANL) facilities. This process will ensure that the material condition of LANL facilities and assets are evaluated, documented, and maintained in order to support safe and reliable plant operations and meet the site long-range planning activities." [26]
AP-MNT-005	Annual Maintenance Work Plan	The Annual Maintenance Work Plan (AMWP) establishes a baseline for a yearly facility maintenance work plan that includes cost, scope, schedule, and resources needed to achieve a maintenance program. This administrative procedure (AP) describes the process for the development of an activity-based AMWP for maintenance work activities to be performed over the forthcoming fiscal year(s) (FY) inclusive of initiatives for maintenance program improvement. [27]
AP-MNT-006	Preventive and Predictive Maintenance	"Implements the requirements and processes for the Los Alamos National Laboratory (LANL) Preventive / Predictive Maintenance (PM/PdM) Program for nuclear and non-nuclear facilities." [28]
AP-MNT-007	Measuring, Analyzing, and Reporting of Maintenance Program Performance	"Establishes the measurement, analysis and reporting process for the Los Alamos National Laboratory (LANL) Maintenance Management Program, which provides the management information necessary to assess maintenance performance and develop action plans for continuous improvement." [29]
AP-MNT-008	Control of Maintenance Tools and Equipment	"Defines the Los Alamos National Laboratory (LANL) maintenance organization's responsibility for the proper procurement, identification, control and handling of maintenance equipment and tools, including calibrated measuring and test equipment (M&TE)." [30]
AP-MNT-010	Maintenance History	"Provides the process for the documentation of Structure, System, and Component (SSC) maintenance history, required in support of the Los Alamos National Laboratory (LANL) maintenance management program. The purpose of the maintenance history program is to document equipment technical data, maintenance performed, and to record system performance. Maintenance history is used to support maintenance activities, trend equipment performance, analyze SSC service life expectancy, and improve equipment reliability." [31]
AP-MNT-013	Deferred Maintenance Identification and Reporting	"Provides guidance and instruction for the identification, recording, maintenance, and reporting of Deferred Maintenance (DM) and Repair Needs (RN) data. Guidance and instructions to manage changes to DM and RN data resulting from corrective maintenance activities at the Los Alamos National Laboratory (LANL) are included." [32]
AP-MNT-014	MSS AP and/or NMMP Exception or Variance Request Process	"Establishes the process for requesting an exception or variance to the Maintenance and Site Services (MSS) APs and/or to the National Nuclear Security Administration (NNSA) approved maintenance program as documented in Los Alamos National Laboratories (LANL) P950, Conduct of Maintenance, procedure." [33]

7.1. LANL maintenance engineering related reports

Table 6 lists relevant reports and documents that were retrieved from a search of the Los Alamos Authors database. The search term "maintenance engineer*" was used in order to capture "engineers" and "engineering." Few reports surfaced – it appears that maintenance engineering has not been a significant focus of the laboratory.

Table 6. List of Los Alamos National Laboratory reports found in the Los Alamos Authors database.

ID	Document Name	Type	Notes
LA-UR-17-24630	Predictive Maintenance Report TA-53 LANSCE	Report	Appears to be an October 2016 inspection report. Lists equipment condition levels for what appears to be various facility equipment at TA-53, uses a three category system: 1) Critical: Examine, Repair, replace ASAP. 2) High: Begin frequent monitoring. Examine, repair, replace during next scheduled shutdown. 3) Low: schedule proactive repair. The report categorizes equipment and shows three "High" pieces. Some equipment is investigated further through a mix of vibration analysis, thermal analysis, and ultrasonic analysis.
LA-UR-19-26373	Fault Detection and Predictive Maintenance for Mechanical Systems	Presentation transcript	The author discusses a student project on predictive maintenance. The author describes using accelerometers to gather data on bearings (such as a bearing in HVAC equipment) This data is used with machine learning to classify bearings into four fault states: 1) Problem detected, 2) Repair & replacement is necessary; 3) bearing is nearing end-of-life; 4) failure is imminent
LA-UR-16-28997	Prioritization of Stockpile Maintenance with Layered Pareto Fronts	Journal article	Uses a Define-Measure-Reduce-Combine-Select (DMRCS) to prioritize maintenance funding. Similar to an analytical hierarchy process, this approach used Pareto Fronts to determine the objectively better choices within a large group of potential options.
LA-UR-19-27403	Predicting and Preventing Failures	Presentation	See notes on LA-UR-19-26373
LA-UR-15-23483	Maintenance Plans for Old/New Equipment	Presentation	Briefly describes maintenance terms and maintenance policies at LANL, then describes the fabrication of a tool used to assist maintenance activities within a CNC glovebox
LAPR-2011-011343	System health assessment	Journal article	Describes System Health Assessment as a concept, and as a means of facilitating prognostics and health management - ways to determine needs of preventive and corrective maintenance and optimize maintenance scheduling.
LA-UR-18-30630	Condition Based Maintenance Implementation for Hydraulic Submarine Actuator Project	Report	Demonstrates a methodology to implement a condition-based maintenance system on submarine components.
LA-UR-18-24167	Unintended Consequences of Bearing Maintenance	Report	Describes a "lessons learned" situation in which maintenance procedures were not adequately followed, resulting in increased troubleshooting and costs after the maintenance took place.
LA-UR-15-24267	Operation & Maintenance Plan for Aboveground Storage Tank Technical Area 16-0980 Generator at the TA-16 Weapons Engineering Tritium Facility (WETF), Los Alamos National Laboratory	Procedure	This is an O&M plan for portions of TA-16. It describes inspection procedures, inspection worksheets / forms, maintenance policies, etc.

8. Industry best-practices

The following case studies have been gleaned from various trade & industry magazines. These case studies demonstrate the implementation of a new or updated maintenance strategy. Each case study provides key takeaways that may be helpful for maintenance managers.

8.1. Case study: BMW automobile manufacturing

The BMW plant in Spartanburg, S.C. has achieved an overall equipment effectiveness of over 90%, with nearly 100% uptime in mission critical areas [34]. BMW plans nearly all maintenance through preventive and predictive strategies. In some cases, this planning occurs up to seven years in advance. BMW only performs scheduled maintenance outside of standard shifts, when the machines are not running as part of the production process. This limits machine downtime. Maintenance forms a core of BMW's manufacturing identity; one board chairman started as a maintenance engineer and seems to drive the company towards predictive and planned maintenance. This top-level focus on maintenance does not make a rigid corporate maintenance structure, however.

The manufacturing plant has four different floor units. Each unit is free to determine their own maintenance strategy. In one case, there is no maintenance department. Rather, each machine operator is trained on and is empowered to make maintenance decisions. A broad "Plant Maintenance Steering Committee," composed of staff from each unit, make larger strategic maintenance decisions. This committee guides maintenance throughout the plant. BMW also has focused on equipment standardization across its business. Finally, BMW heavily uses its CMMS and has transitioned to a condition-based maintenance approach with sensors throughout its assembly lines.

A key point from this case study is: train employees to rely on the CMMS and empower them to be owners of maintenance. By accomplishing this, preventive and predictive maintenance will become readily achievable.

8.2. Case study: Alcoa aluminum foundry

Alcoa operates a primary metals facility in southern Indiana, originally constructed in the 1950s [35]. This Alcoa plant was attempting to be cost-competitive despite the outdated facility. A strong maintenance program was a key factor in cost reductions. Compared to peers, maintenance was roughly 40% more costly at Alcoa. In order to reduce costs, Alcoa implemented a "Reliability Excellence Process" that aimed to define a formal partnership between the maintenance and operations divisions. This involved subject matter expert interviews to understand where the old partnership broke down. The program developed a plan to close these gaps through new policies and employee empowerment.

Here, management commitment was key to instituting cultural shifts. Management focused on improving overall equipment effectiveness. There were definitional issues with their KPIs – their "planned and scheduled work" KPI also included scheduled corrective maintenance tasks,

misleading leadership into thinking they had good preventive maintenance scheduling. Alcoa hired more planners to help with maintenance planning, and settled on a ratio of 20-1 operators to maintenance planners (down from 35-1). They also created a dedicated staging area to house spare parts in order to track inventory and reduce lead times.

A key point from this case study is updating institutional policies to address problems-of-the-day. Alcoa additionally invested in maintenance managers, who developed efficiency solutions such as parts tracking through a more-utilized CMMS.

8.3. Case study: Texas Instruments wafer production

The Texas Instruments' (TI) wafer production plant in Dallas, Texas was written about in [36]. The facility produces various semiconductors across four neighboring factories. Maintenance responsibilities are shared between a facilities group (responsible for building envelope, but also equipment installation) and an equipment engineering group (responsible for equipment post-installation).

The TI facility shifted towards a zero waste goal, including zero major interruptions (caused by corrective maintenance). To achieve this goal, the staff focused on traditional reliability metrics but also performed a strategic reanalysis. This analysis reconsidered maintenance policies and the maintenance concept, using a framework centered on the question "if we were to restart from scratch, what would we do differently?" Maintenance managers also consulted other facilities and operations across the company, including production plants in other countries, in order to share maintenance knowledge and best-practices.

This analysis uncovered four points of inefficiency. 1) There was a history of over-maintenance of non-critical components; 2) not enough effort was focused on critical components; 3) there were too many unique assets to cover; 4) available worker time had declined because of retirements and staff shuffling.

TI developed a new maintenance concept, where finite resources were directed at only the most critical components and equipment. Equipment deemed critical received continuous condition monitoring for predictive maintenance. Lower level equipment received intermittent condition monitoring through inspection, or schedule-prescribed preventive maintenance. Other non-critical equipment were allowed to run to failure. In this way, cost and impact were balanced and optimized through a risk analysis approach.

A key point from this case study is: maintenance prioritization is incredibly beneficial within the context of limited budget and limited staff.

9. Establishing a maintenance program

The trade journal "Maintenance & Engineering" recently published an article titled *Seven Steps to Maintenance Strategy* that describes how to establish a maintenance program [37]. The article claims that maintenance is often constrained to a task list, whereas a robust maintenance strategy creates focus and ambition for the maintenance team and relevant stakeholders. Each of the seven steps to create this strategy is described below.

9.1. Understand current maintenance maturity

The overall maturity of a maintenance program ranges from only corrective maintenance actions to a full predictive maintenance environment. As maturity increases, maintenance staff and operators gain needed skills and experience to progress the maintenance program. This progression is desired, but takes time to develop. Understanding an organization's current maturity is critical to implementing effective change and effective maintenance policies – in order to benefit the organization immediately, but also to advance the maintenance program.

The levels of maturity, and best-practices at those levels, are provided in Table 7.

Table 7. Maintenance program maturity levels, adapted from [37].

Level of maintenance program maturity	Best Practices
Breakdown maintenance	Reduce lead times, create effective technician training
Planned (corrective) maintenance	Identify problematic equipment Identify needed spares and tooling Optimize maintenance times
Preventive maintenance	Prioritize critical machines Define failure modes Optimize preventive maintenance cycles
Basic predictive maintenance	Replace preventive methodologies with condition-based methodologies Optimize labor planning
Connected predictive maintenance	Develop real-time condition monitoring Integrate data sources Optimize CMMS

9.2. Develop and track KPIs

Key performance indicators (KPIs) are critical for a more effective maintenance program. Maintenance managers ought to determine and develop the means to track important KPIs for the organization. These may include metrics such as machine or system availability, overall equipment effectiveness, or costs and time used.

9.3. Deploy an effective asset strategy

Next, maintenance managers ought to determine which maintenance policies should be implemented for their equipment. Typically, policies fall into two or three general categories (see Figure 6), corrective maintenance ("run to failure"), preventive maintenance, and predictive maintenance.

Preventive maintenance is typically based on manufacturer recommendations, sometimes provided in an operations or service manual. If an organization lacks failure data, it is best to follow these recommendations with the understanding that they may be somewhat overly-conservative. Managers may adjust these intervals as they gain experience and failure data with the equipment.[17]

Over time, these asset strategies may to shift to minimize corrective maintenance (unless run-to-failure is genuinely cheaper than a repair strategy). Maturing programs will also shift towards predictive maintenance.

9.4. Involve all stakeholders

As seen in the case studies provided in Section 8, stakeholder involvement is critical to develop a maintenance culture. Equipment operators can be trained to provide formal and informal machine inspections, and to provide this data to the CMMS. Maintenance planners and technicians ought to have a close working relationship to develop and understand maintenance actions. Finally, in some cases equipment designers should be involved to receive feedback on maintenance actions and frequency.

9.5. Communicate economic benefits

Effective maintenance strategies meet the needs and objectives of upper management. Often, this involves reducing costs or providing an economic benefit. Maintenance is often seen as a cost area rather than a division or group that provides a direct economic benefit. These benefits should be quantified and used to better communicate the effects of a maintenance strategy. Some of these benefits include increased availability, extended equipment lifetimes, and/or mitigating safety or quality risks.

9.6. Build maintenance capability

Over time, a maintenance group should be building capability as the maintenance program matures. This involves investing in training and data analysis. Maintenance planners should be provided time to investigate datasets within a CMMS in order to better plan maintenance actions. Other capabilities to improve include knowledge & procedure management, maintenance preparation & execution, and cost estimating.

---

[17] During a repair activity, the following data ought to be collected:
- Age of item at failure
- Usage of item
- Reason for failure
- Usage mode, intensity, operating environment
- Symptoms prior to failure
- Actions take to rectify failure
- Repair data including parts used, time to repair, etc.

9.7. Prompt culture change

Through the above steps and tasks, a maintenance manager may begin shifting the broader organizational culture. Typically, a shift is seen from focusing on short-term budgets to the longer-term benefits of maintenance planning and predictive maintenance actions.

Works Cited

[1]   M. Ben-Daya, U. Kumar and D. P. Murthy, "Preface," in *Introduction to Maintenance Engineering*, Hoboken, John Wiley & Sons, Inc., 2016, p. xxxii.

[2]   "Skywise," Palantir, [Online]. Available: https://www.palantir.com/solutions/skywise/. [Accessed 1 April 2021].

[3]   "IBM Maximo Application Suite for Predictive Maintenance," IBM, [Online]. Available: https://www.ibm.com/products/maximo/predictive-maintenance. [Accessed 1 April 2021].

[4]   M. Ben-Daya, U. Kumar and D. P. Murthy, "System Degradation and Failure," in *Introduction to Maintenance Engineering*, Hoboken, John Wiley & Sons, Inc., 2016, pp. 51-71.

[5]   M. Ben-Daya, U. Kumar and D. P. Murthy, "Basics of Reliability Theory," in *Introduction to Maintenance Engineering*, Hoboken, John Wiley & Sons, Inc., 2016, pp. 25-50.

[6]   P. A. Gagniuc, Markov Chains: From Theory to Implementation and Experimentation, First Edition, Hoboken: John Wiley & Sons, Inc, 2017.

[7]   MIL-STD-882E, "Department of Defense Standard Practice: System Safety," Department of Defense, 2012.

[8]   DOE G 413.3-7A, "Risk Management Guide," Department of Energy, 2015.

[9]   C. S. Wasson, "System Reliability, Maintainability, and Availability," in *System Engineering Analysis, Design, and Development*, Hoboken, John Wiley & Sons, Inc., 2015, pp. 721-791.

[10] G.-A. Klutke, P. C. Kiessler and M. A. Wortman, "A Critical Look at the Bathtub Curve," *IEEE Transactions on Reliability,* vol. 52, no. 1, 2003.

[11] W. Nelson, Accelerated Testing: Statistical Models, Test Plans, and Data Analyses, New York: Wiley, 1990.

[12] A. M. Smith, Reliability-Centered Maintenance, New York: McGraw-Hill, 1992.

[13] J. C. Lee and N. J. McCormick, "Reliability of Multiple-Component Systems," in *Risk and Safety Analysis of Nuclear Systems*, Hoboken, John Wiley & Sons, Inc., 2011, pp. 85-104.

[14] P. D. T. O'Connor and A. Kleyner, "Life Data Analysis and Probability Plotting," in *Practical Reliability Engineering*, West Sussex, United Kingdom, John Wiley & Sons, Ltd., 2012, pp. 70-107.

[15] D. F. Percy, "Maintenance Based on Limited Data," in *Complex System Maintenance Handbook*, London, Springer, 2008, pp. 133-154.

[16] W. M. Bolstad, Introduction to Bayesian Statistics, Hoboken: John Wiley & Sons, Inc., 2007.

[17] M. Ben-Daya, U. Kumar and D. P. Murthy, "Maintainability and Availability," in *Introduction to Maintenance Engineering*, Hoboken, John Wiley & Sons, Inc., 2016, pp. 155-169.

[18] M. Ben-Daya, U. Kumar and D. P. Murthy, "Maintenance - Basic Concepts," in *Introduction to Maintenance Engineering*, Hoboken, John Wiley & Sons, Inc., 2016, pp. 73-105.

[19] S. Nakajima, Introduction to Total Productive Maintenance, Productivity Press, Inc., 1988.

[21] J. Trout, "What is Overall Equipment Effectiveness (OEE)?," Reliable Plant, [Online]. Available: https://www.reliableplant.com/Read/11785/overall-equipment-effectiveness. [Accessed 16 3 2021].

[22] M. Ben-Daya, U. Kumar and D. P. Murthy, "Optimal Maintenance," in *Introduction to Maintenance Engineering*, Hoboken, John Wiley & Sons, Inc., 2016, pp. 301-309.

[23] Los Alamos National Laboratory, "P950 Conduct of Maintenance," 2020.

[24] Los Alamos National Laboratory, "AP-MNT-002 Rev 5 Seasonal Facility Preservation," 2019.

[25] Los Alamos National Laboratory, "AP-MNT-003 R.5 Determining Maintenance Facility, Equipment and Tool Needs," 2019.

[26] Los Alamos National Laboratory, "AP-MNT-004 R3 Facility Condition Inspection," 2019.

[27] Los Alamos National Laboratory, "AP-MNT-005 R9 Annual Maintenance Work Plan," 2020.

[28] Los Alamos National Laboratory, "AP-MNT-006, R6 Preventive and Predictive Maintenance," 2020.

[29] Los Alamos National Laboratory, "AP-MNT-007 R.7 Measuring, Analyzing, and Reporting of Maintenance Program Performance," 2020.

[30] Los Alamos National Laboratory, "AP-MNT-008, R4 Control of Maintenance Tools and Equipment," 2011.

[31] Los Alamos National Laboratory, "AP-MNT-010 R7 Maintenance History," 2020.

[32] Los Alamos National Laboratory, "AP-MNT-013 R3 Deferred Maintenance & Repair Needs Identification and Reporting," 2019.

[33] Los Alamos National Laboratory, "AP-MNT-014, R1 MSS AP and/or NMMP Exception or Variance Request Process," 2019.

[34] P. V. Arnold, "World-class OEE Drives BMW Plant," Reliable Plant, [Online]. Available: https://www.reliableplant.com/Read/5197/bmw-reliability. [Accessed 21 4 2021].

[35] P. V. Arnold, "Alcoa Pursues Reliability Excellence," Reliable Plant, [Online]. Available: https://www.reliableplant.com/Read/256/alcoa-reliability. [Accessed 21 4 2021].

[36] P. V. Arnold, "Maintenance and Lubrication Success at Texas Instruments," Machinery Lubrication, 7 2009. [Online]. Available: https://www.machinerylubrication.com/Read/2304/texas-instruments-reliability. [Accessed 23 4 2021].

[37] P. Adams, "Seven steps to maintenance strategy," Maintenance & Engineering, 9 8 2019. [Online]. Available: https://www.maintenanceandengineering.com/2019/08/09/seven-steps-to-maintenance-strategy/. [Accessed 20 4 2021].

[38] MIL-HDBK-470A, "Designing and Developing Maintainable Products and Systems," U.S. Department of Defense, 1997.

[39] M. Ben-Daya, U. Kumar and D. P. Murthy, "Maintenance Management," in *Introduction to Maintenance Engineering*, Hoboken, John Wiley & Sons, Inc., 2016, pp. 391-419.

[42] J. Moubray, Reliability-Centered Maintenance, New York: Industrial Press, Inc., 1991.